Building Internal Clouds: Practical Guidance for a Difficult Challenge

Chris Wolf

Research Vice President

chris.wolf@gartner.com

Twitter: @cswolf

Gartner

Building Internal Clouds

Agenda

- Burton Group HIaaS Reference Architecture
- Building an Internal Cloud

HIaaS Reference Architecture

Cloud HIaaS Reference Architecture

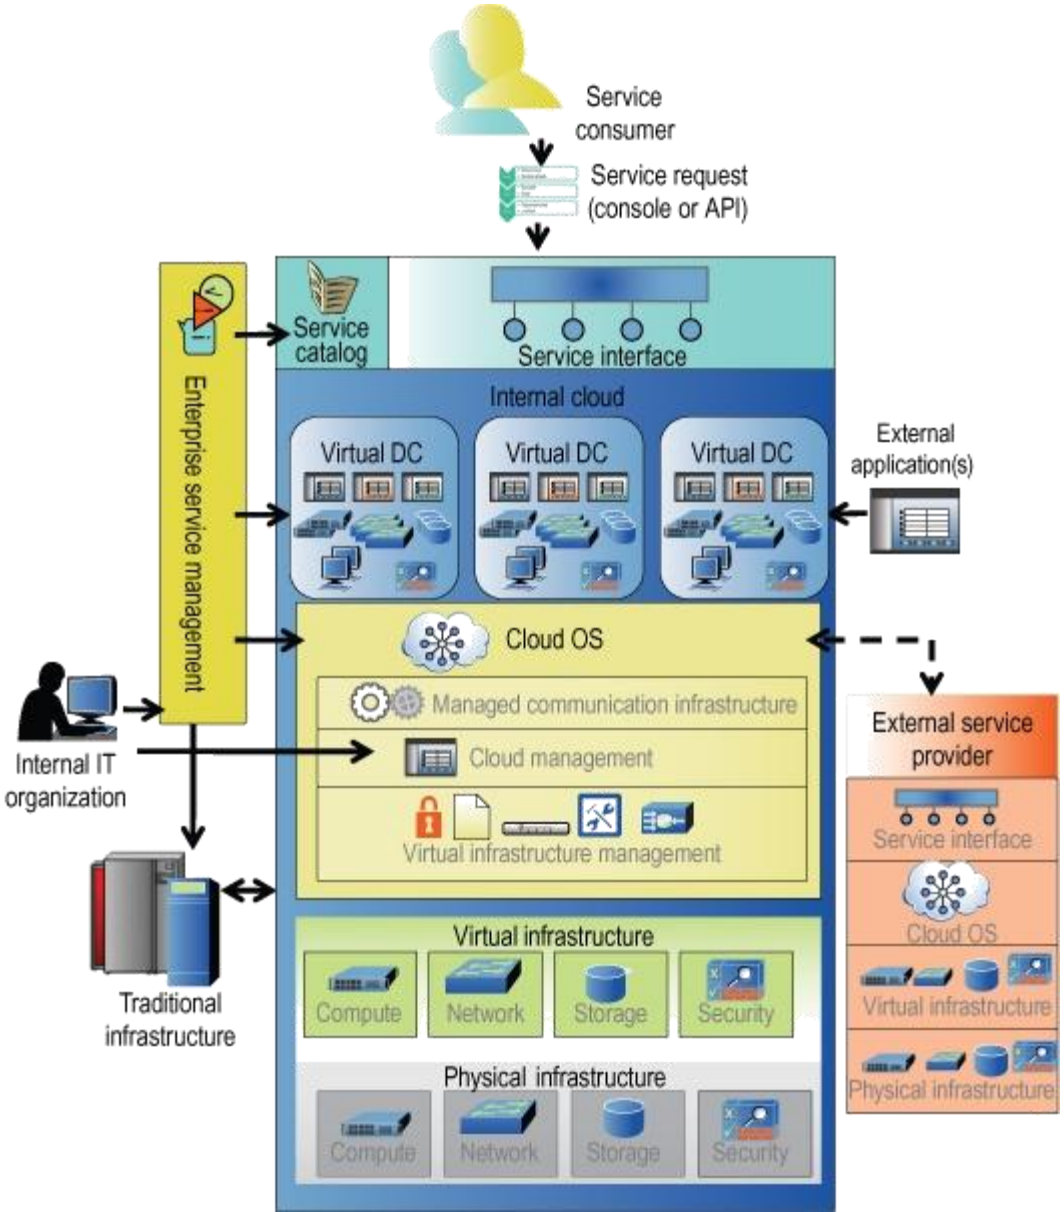

Gartner

Cloud HIaaS Reference Architecture

Virtual Data Center

- Secure virtual container that encapsulates a service or application
- Defined by standard metadata format (e.g., OVF)
- Contents:
  - 1 or more VMs
  - Managed distributed virtual switches
  - Routers, firewalls, WAN accelerators
  - Application delivery controllers
  - Intrusion prevention systems (IPS) and IDS
  - Storage virtualization appliances

Building the Internal Cloud

Internal Cloud Maturity Model

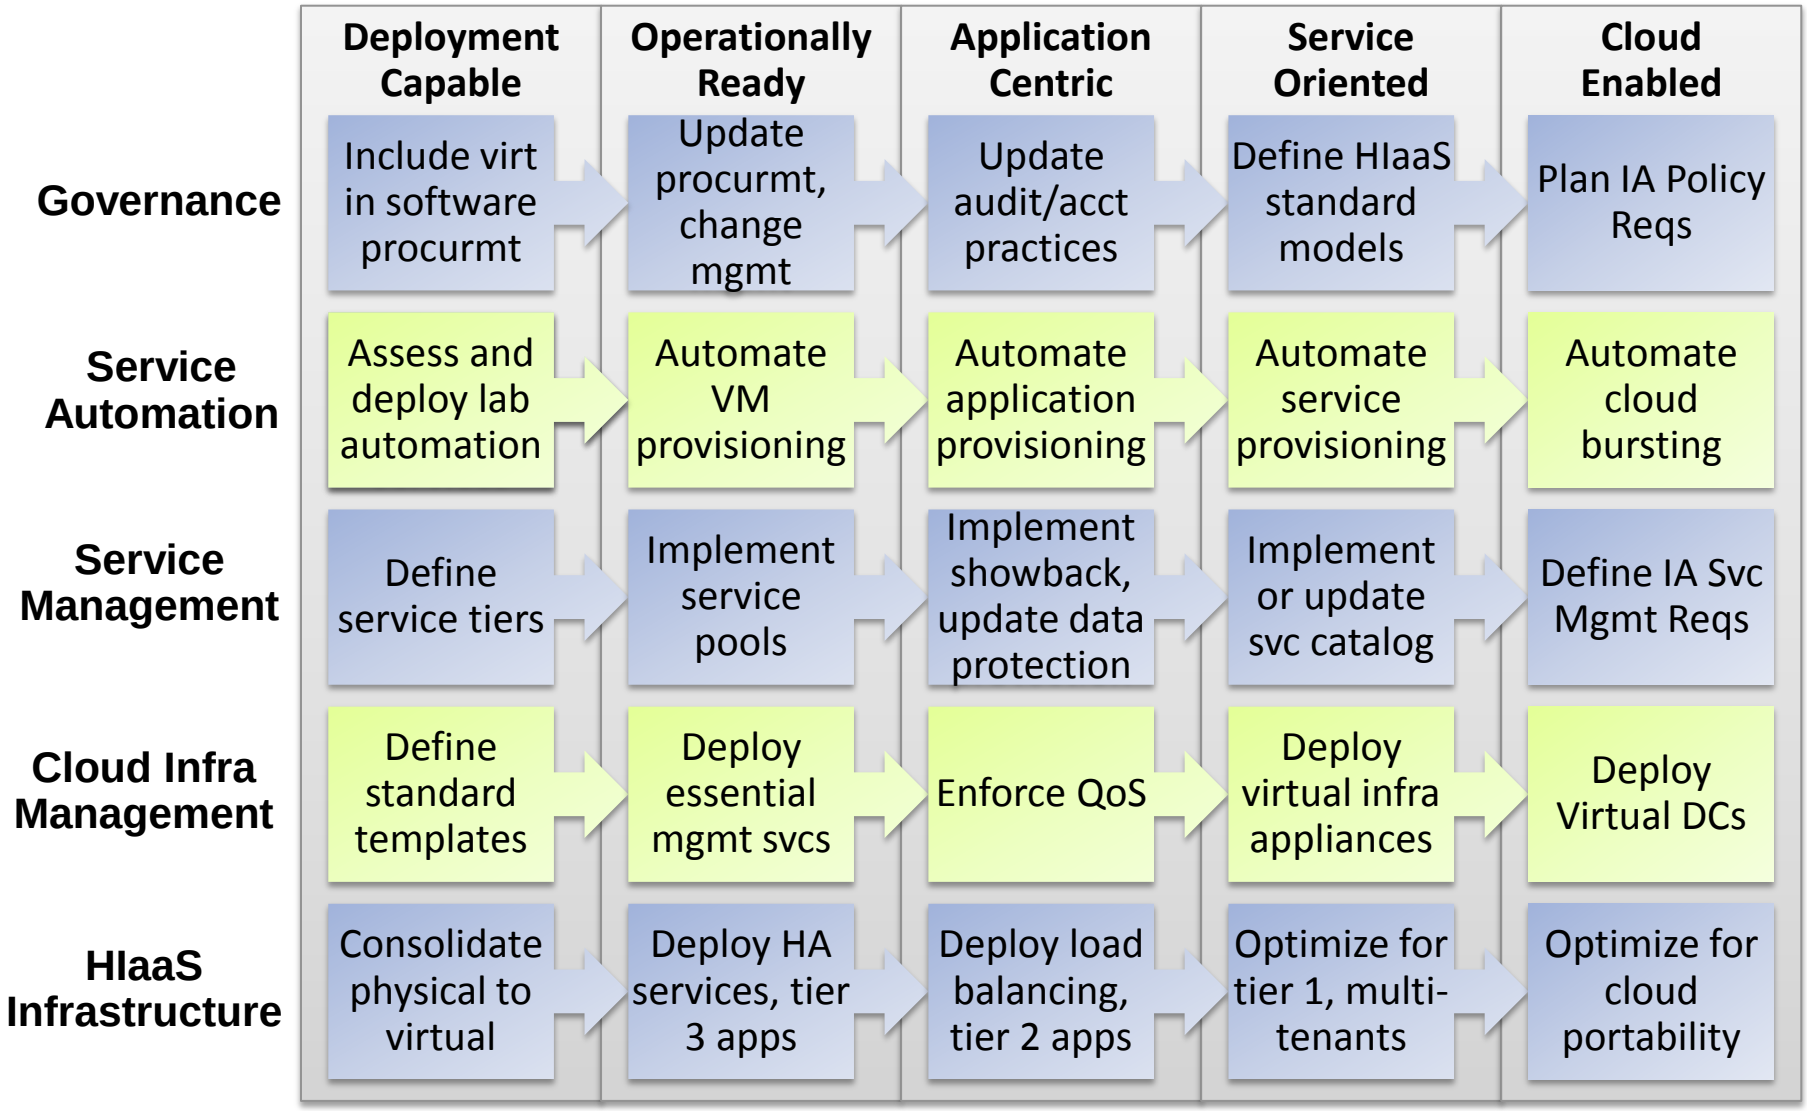

Internal Cloud Pre-work

- Devise a cloud computing strategy
- Map security, application, identity, and information management to cloud strategy
- Assess and deploy the x86 server virtualization infrastructure
- Assess and deploy storage infrastructure
- Assess and deploy network infrastructure
- Assess and deploy compute infrastructure
- Deploy and manage non-critical virtual workloads

Internal Cloud Pre-work

Devise a Cloud Computing Strategy

- Drives infrastructure, management, and service automation requirements
- Definitions:
  - ***Public cloud:*** IT capability as a service that providers offer to consumers via the public Internet
  - ***Private cloud:*** IT capability as a service that providers offer to a select group of customers.
  - ***Internal cloud:*** IT capability as a service that an IT organization to its own business (subset of private)
  - ***External cloud***: An IT capability as a service offered to a business that is not hosted by its own IT organization.
  - **Hybrid cloud**: IT capabilities offered as a service that spans internal and external clouds

Internal Cloud Pre-work

Cloud Tiered Architecture

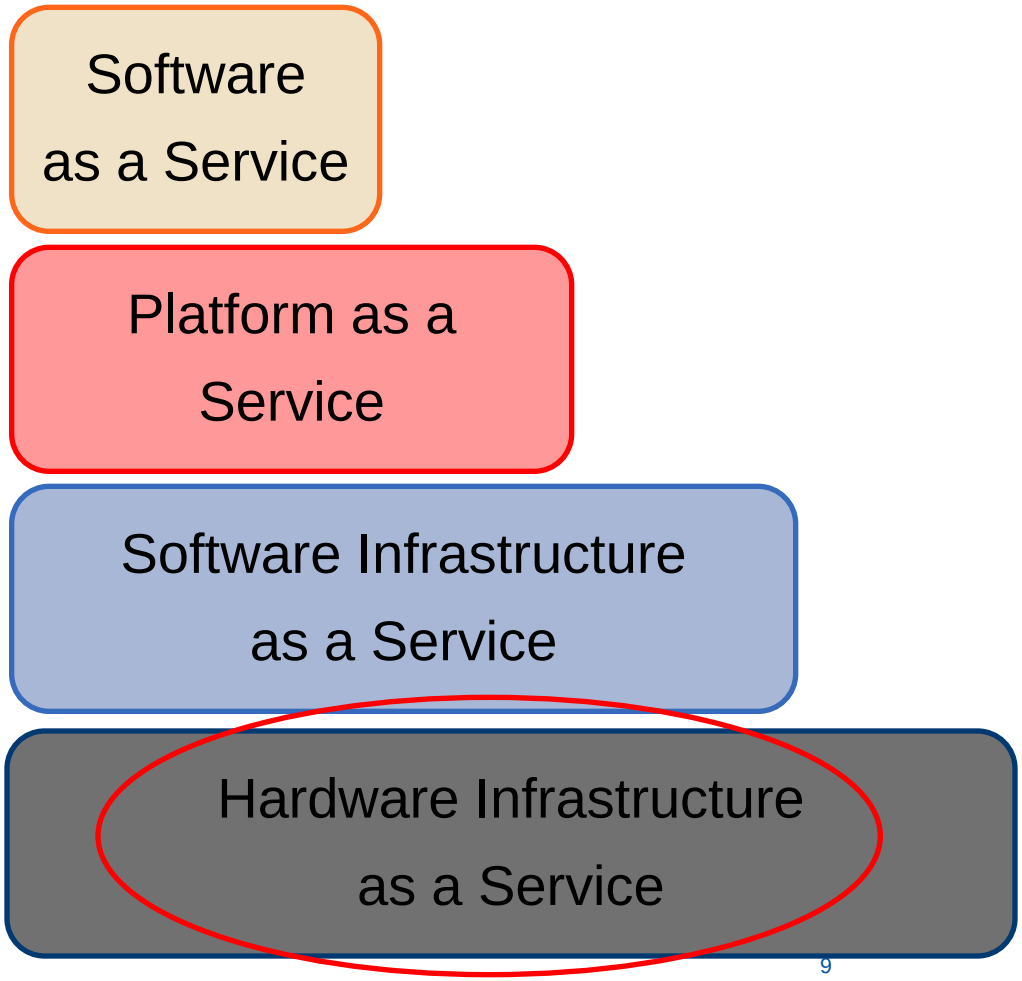

Examples
Google Apps, Salesforce.com, backup as a Service
Microsoft Azure, Force.com, Google App Engine, Oracle's Cloud Strategy
Data services (e.g., MS SQL DB, Amazon SimpleDB), content distribution, Identity and security providers
EC2, system hosting providers (e.g., BT, AT&T, Sprint), virtualization vendors

Internal Cloud Pre-work

Strategy: Infrastructure Transparency

- A few quotes:
  - "All we should care about is the app and the SLA"
  - "If the provider gives me the virtual resources I need, why should I care about the underlying hardware?"
- Why we need it
  - Low-level features impact performance and security
    - Intel Extended Page Tables (EPT) and Trusted Execution Technology (TXT)
    - AMD Rapid Virtualization Indexing (RVI)
  - Security policy and zoning enforcement
  - Reliability: organizations willing to pay more for cloud infrastructure running on brands they trust

Internal Cloud Pre-work

Devise a Cloud Computing Strategy: Resources

- Building a Solid Cloud Adoption Strategy: Success by Design
- Cloud Computing Tiered Architecture
- Planning Considerations for Externalization and Cloud Computing
- Developing a Cloud Computing Security Strategy

Internal Cloud Pre-work

Map security, application, identity, and information management to cloud strategy

- Close alignment is required to automate infrastructure operations while meeting service level and compliance requirements

Internal Cloud Pre-work

Assess and deploy the x86 server virtualization infrastructure

- The hypervisor is core infrastructure technology
- Replacing the hypervisor is costly:
  - Downtime
  - Training
  - Re-architecting IT and business processes

Internal Cloud Pre-work

Assessing Virtualization Solutions

- Gartner Magic Quadrant for x86 Server Virtualization Infrastructure

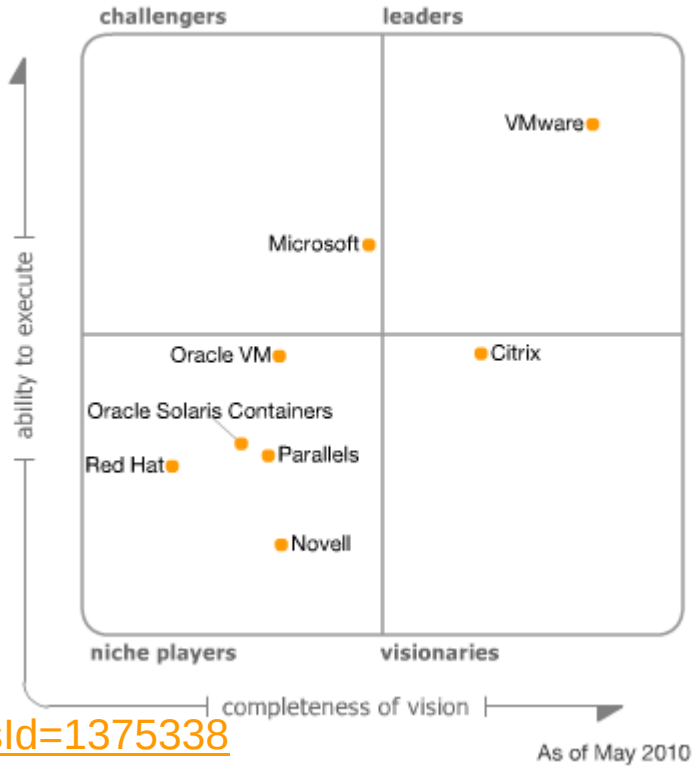

http://www.gartner.com/resId=1375338

Gartner

Internal Cloud Pre-work

Assessing Virtualization Solutions

- Burton Group Evaluation Criteria and Product Scorecards
  - A product must meet 100% of *required* criteria to be deemed *enterprise-ready*

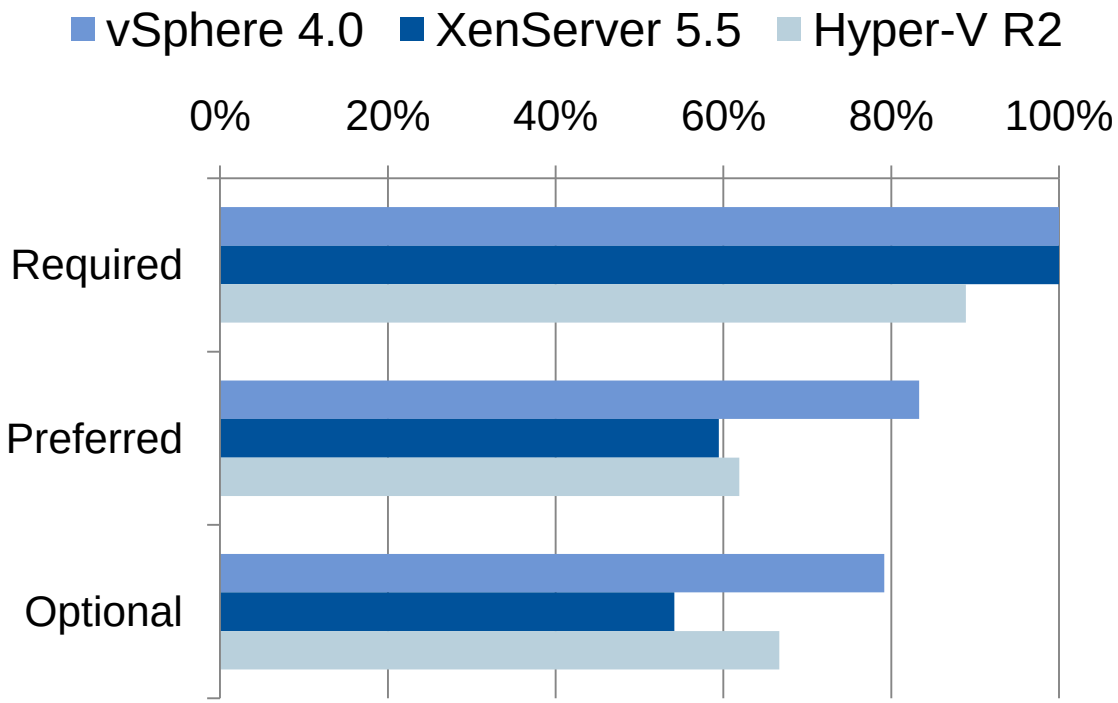

Gartner

Internal Cloud Pre-work

Solution Assessment Resources

- *Magic Quadrant for x86 Server Virtualization Infrastructure*
  http://www.gartner.com/resId=1375338
- *Server Virtualization Hypervisors*
  http://www.burtongroup.com/Client/Research/Document.aspx?cid=1569
- *VMware vSphere 4.0*
  http://www.burtongroup.com/Client/Research/Document.aspx?cid=1653
- *Citrix XenServer 5.5*
  http://www.burtongroup.com/Client/Research/Document.aspx?cid=1756
- *Windows Server 2008 R2 Hyper-V*
  http://www.burtongroup.com/Client/Research/Document.aspx?cid=1654

Internal Cloud Pre-work

Single vs. Multi Vendor

- The Question is Not Single vs. Multi-hypervisor
- …but what percentage of hypervisors are managed by IT?
- Practically all organizations have at least 2 of the following:
  - VMware: ESX, Player, Workstation, Fusion, ACE
  - Microsoft: Hyper-V, Virtual Server, Virtual PC
  - Citrix: XenServer, XenClient
  - Oracle: OVM, Sun VirtualBox
  - Parallels: Workstation, Server
  - Open source Xen and KVM

Internal Cloud Pre-work

We're trying to standardize on server hardware to reduce capacity, configuration, and device management costs

- Single vendor:
  - Fewer tools, fewer images
  - Cloud portability
  - Enforceable standards
  - Higher lock-in
- Multi-vendor
  - More tools and images
  - Portability more complex (VM conversion, vendor support, internal certification)
  - Reduced lock-in and acquisition costs

Internal Cloud Pre-work

Common Multi-Vendor Scenarios

- Separate server and desktop business units
  - VMware or Hyper-V for servers
  - XenServer for XenApp farm and virtual desktops
- Separate sites
- Non-critical workloads (e.g., training, dev, test)
  - Leverage free hypervisor (XenServer, Microsoft Hyper-V, ESXi, Xen, or KVM)

Internal Cloud Pre-work

Assess and Deploy Storage Infrastructure

- Storage topology: FC/NFS/iSCSI/FCoE
- Array features:
  - Thin provisioning
  - Deduplication
  - Integration with virtualization management
    - Automated provisioning
    - Snapshots
    - vStorage APIs for Array Integration (VAAI)

Internal Cloud Pre-work

Assess and Deploy Storage Infrastructure

- Additional Information:
  - Storage for Server Virtualization
  - Storage Ecosystem
  - Selecting an Enterprise-Ready Disk Array
  - Enterprise Disk Array Selection Methodology
  - It’s the Data That Matters: Storage Considerations for Data Center Consolidation
  - Efficient Storage: Well Within Reach

Internal Cloud Pre-work

Assess and Deploy Network Infrastructure

- Burton Group position: network (not storage) admins should manage virtual networks
- Distributed virtual switch recommended
- Most organizations physically isolate security zones at early stages of maturity
- Additional information:
  - "Networks for Dynamic Virtualized Data Centers" (overview)
  - "NTS Root Template" (reference architecture template)

Security Zone Segmentation

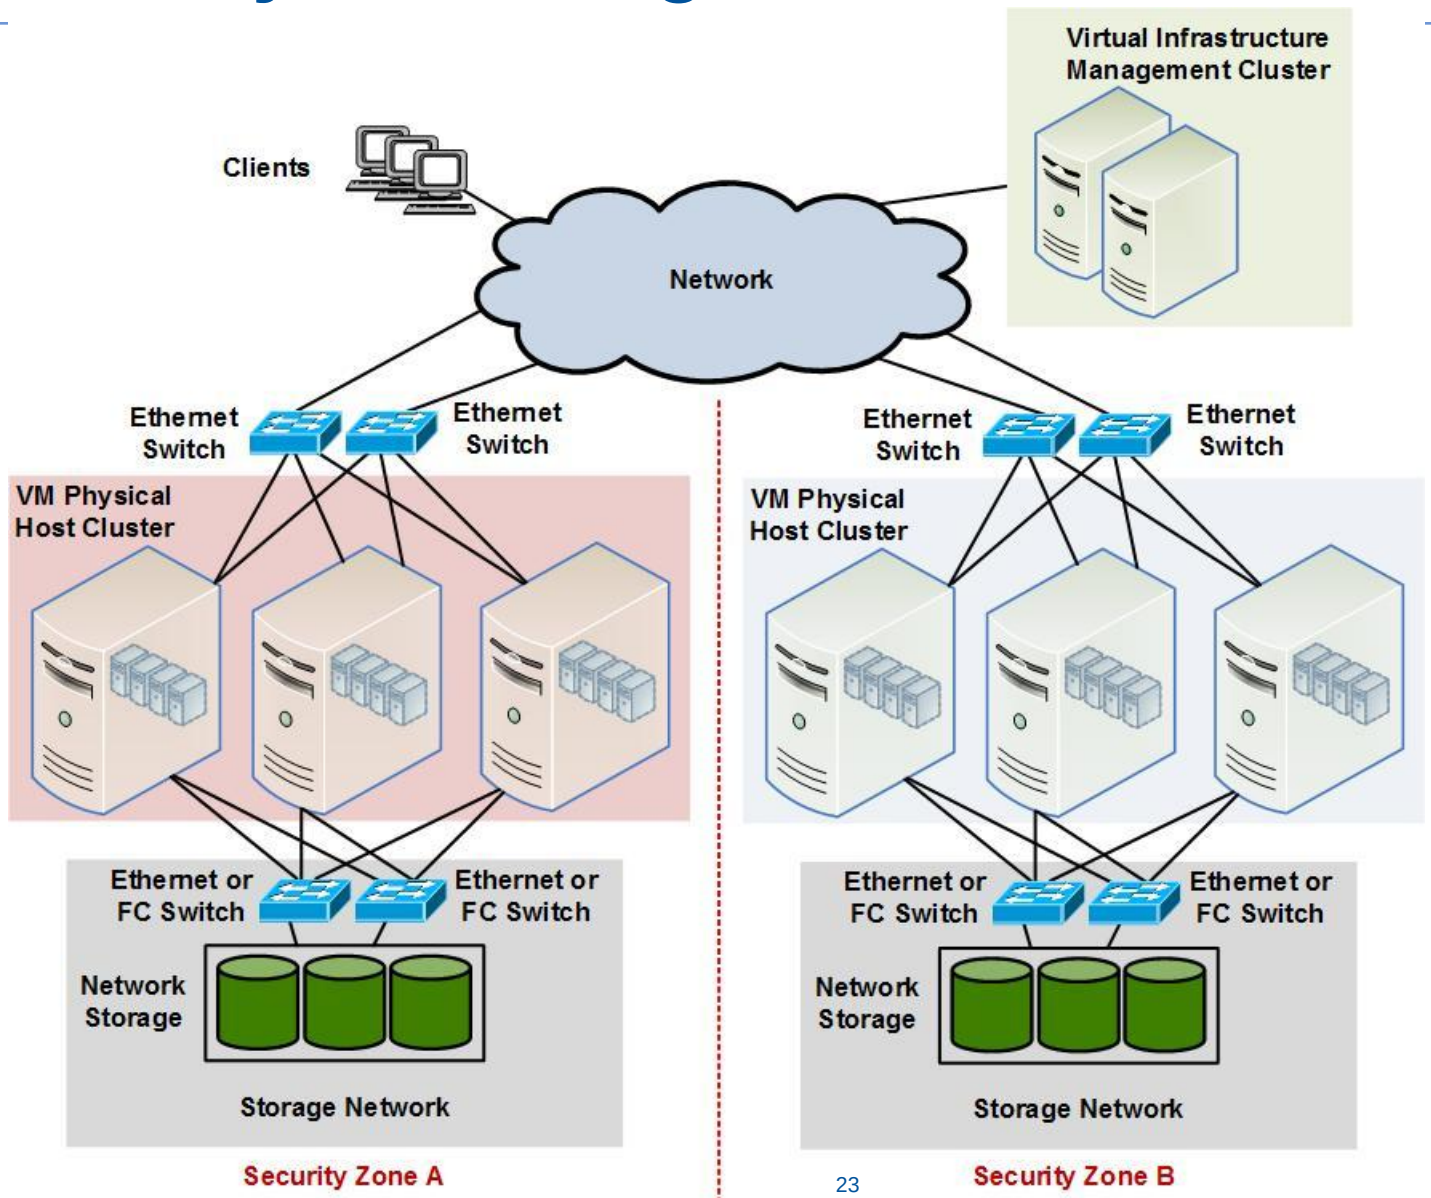

Gartner

Internal Cloud Pre-work

Assess and deploy compute infrastructure

- Second generation platforms that support hardware-assisted memory virtualization (i.e., Intel EPT and AMD RVI) highly recommended
- Processor consistency impacts mobility and live migration functionality
- Additional information:
  - "Selecting the Correct Server Platform" (decision point)
  - "The Commodity Server is Being Reshaped by Server Virtualization" (overview)
  - "x86 Blade Server Platforms: Cutting Through the Hype"(report)

Pick the Right Hardware

- Blade density looks great on paper…
  - Blades vs. rack mount in a standard rack

	Typical blade	2U 4-socket server	Ratio
Processors/rack	128 (4-sockets x 32 blades)	84 (21 x 4-socket servers)	**65%**
Memory/rack	4 TB (128 GB x 32 blades)	5.376 TB (21 servers x 256 GB)	**134%**
I/O bandwidth	2.048 Tbit/s	4.368 Tbit/s	**213%**
Disks	64 x 2.5” (2 drives/blade)	168 x 2.5” (8 drives x 21 servers)	**262%**
Power	19.2 KW	13.650 KW	**71%**

  - Rack mount has more memory, storage & I/O throughput available
  - The only advantage for blades is the number of processors

Server Platform Comparison

Server	Max Memory	CPU	Expansion
Cisco UCS B-200 M1	96 GB (12 slots)	2 quad core Xeon 5500s	1 Mezzanine adapter
Cisco UCS B-250 M1	384 GB	2 quad core Xeon 5500s	2 Mezzanine adapters
Dell PowerEdge r710	144 GB (18 slots)	2 quad core Xeon 5500s	2 PCIe x8, 2 PCIe x4
HP Proliant DL380 G6	144 GB (18 slots)	2 quad core Xeon 5500s	Up to 6 (PCIe, PCI-X)
IBM x3650 M2	128 GB (16 slots)	2 quad core Xeon 5500s	4 PCI-Express

Deployment Capable

- Key characteristics:
  - Non-critical dev/test/training servers are virtualized
  - IT administrators and developers are able to provision test systems using a self-service web interface
  - Standard VM templates and service tiers are defined
- Objectives:
  - Consolidate physical servers onto VMs
  - Define standard VM templates
  - Define VM service tiers
  - Assess and deploy a lab automation platform
  - Deploy/manage non-critical virtual workloads

Deployment Capable

Consolidate physical servers onto VMs

- Reduce vCPU requirements when possible
- Split multi-purpose servers into single function VMs
- Use free tools (e.g., VMware Converter) for basic P2V conversions
- Additional information:
  - P2V: Migrate or Migraine

Deployment Capable

Define standard VM templates

- Template per guest OS type and number of vCPUs
  - 1 template for uniprocessor and separate template for SMP
- Storage and VM guest OS settings are important too
- Not just about the VMs:
  - Server virtualization hypervisor physical hosts
  - Shared or dedicated storage
  - Virtual and physical networks
  - Security

VM Guest Storage Optimization

- Use multiple virtual disks to segregate data
  - Paging file, temp files on one disk
  - OS files on one disk
  - Data files, application files on one or more disks
- Improves image backup and asynchronous replication efficiency

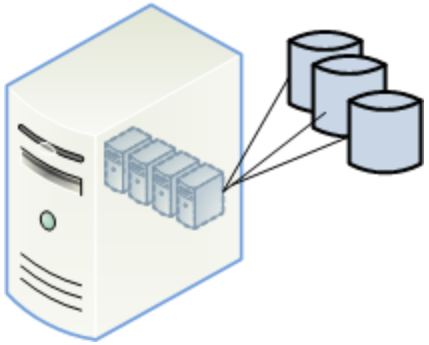

Paging, Temp files
OS files
Data, Application files

Guest Storage Optimization

Going a Step Further

- Separate replication schedules per data tier
- Leave page/swap/temp behind? Impact on SRM?
  - http://media.netapp.com/documents/tr-3671.pdf

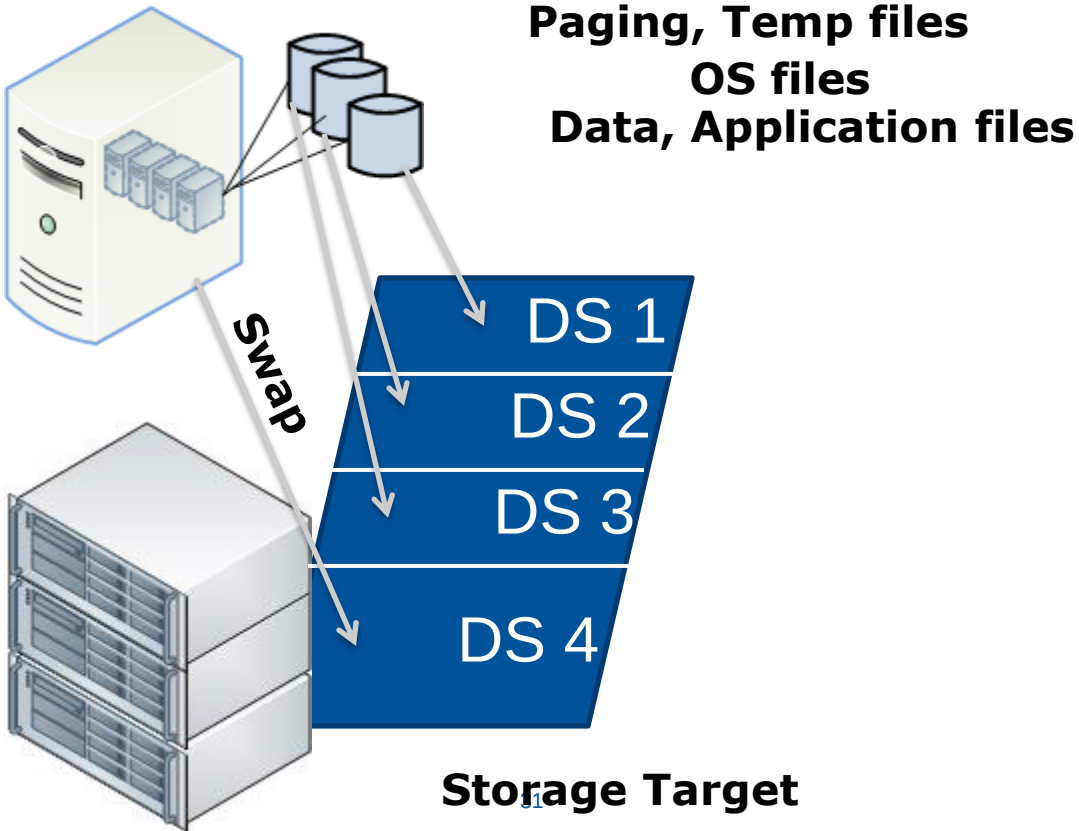

Gartner

Deployment Capable

Define standard service tiers

- 3 to 12 standard service tiers
  - Defines availability, performance, security, and data protection requirements
- Examples:
  - Level 1: No high availability, no data protection, multi-tenant shared infrastructure
  - Level 2: No high availability, image data protection, multi-tenant shared infrastructure
  - Level 3: High availability, image data protection, multi-tenant shared infrastructure

Deployment Capable

Service tier definition attributes

- **Relative priority:** Priority assigned to VM or service; used to prioritize resource access during periods of contention
- **CPU:** number of vCPUs, relative priority, and reserved compute requirements
- **Memory:** amount of VM memory (e.g., 1GB or 8GB) and reserved memory requirements
- **Availability:** VM or service availability requirements, such as no availability, high availability failover, or fault tolerance
- **Recovery time objective (RTO):** Allowable time for VM to be recovered following an outage (e.g., OS, server, cluster, or site failure)
- **Recovery point objective (RPO):** Determines the maximum allowable data loss (e.g., 15 minutes, 2 hours, 24 hours, or 7 days)
- **Security zoning/isolation requirements:** Identifies data and application processing zoning requirements (e.g., full physical isolation, dedicated network and storage, or allows multitenancy but requires unique VLAN membership or virtual switch access control lists [ACLs])

Deployment Capable

Assess and deploy a lab automation platform

- First level of automation for HIaaS
- Common characteristics:
  - Provide self-service provisioning web interface, allowing users to provision one or more VMs into a lab environment
  - Allow users to clone and share a test environment with other users
  - Allows administrators to define a time-to-live (TTL) for lab VMs, enabling automatic system de-provisioning at a pre-determined time
  - Include a quota management system to limit the number of VMs that a given user can have deployed at a given time

Deployment Capable

Include Virtualization Considerations in Procurement Processes

- Virtualization platform support a requirement in RFPs
- Negotiate support policies for explicit clarification on how virtualization is supported by vendor

Operationally Ready

- Key characteristics:
  - The organization is comfortable with the tools and processes associated with migrating physical systems to virtual machines
  - Tier 3 (i.e., low priority or non-critical) services and applications are virtualized
  - IT administrators are able to provision production VMs using a self-service web interface
  - Essential management services (e.g., capacity, configuration, and lifecycle management) are deployed
  - Service pools are deployed
  - The underlying virtual and physical infrastructure is highly available

Operationally Ready

- Objectives:
  - Deploy high availability and virtualize tier 3 applications and services
  - Deploy essential management services
  - Implement service pools
  - Automate VM provisioning
  - Update procurement and change management processes

Operationally Ready

Deploy essential management services

- Capacity management
- Configuration management
- Lifecycle management
- Application diagnostics
- Orchestration/service automation

Application Centric

- Key characteristics:
  - Tier 2 services and applications are virtualized
  - Users are able to provision applications using a self-service web interface
  - A chargeback tool has been deployed so that the IT organization can provide “showback”

Application Centric

- Objectives:
  - Deploy load balancing and virtualize tier 2 applications and services
  - Enforce quality of service (QoS)
  - Implement showback and update data protection
  - Automate application provisioning
  - Update audit and accounting processes

Standard Security Models

- Examples:
  - **Level A:** Dedicated physical and virtual infrastructure, including dedicated server and networked storage assets.
  - **Level B:** Dedicated virtual and physical server infrastructure, shared/logically zoned storage infrastructure (clients receive dedicated LUNs, but data traverses a shared physical SAN).
  - **Level C:** Shared virtual and physical infrastructure, isolation provided by dedicated virtual security appliances (e.g., VM firewalls, IDS, IPS).
  - **Level D:** Shared virtual and physical infrastructure, no appliance-based segmentation and isolation (isolation provided via VLANs).
  - Follow the Cloud Security Alliance: www.cloudsecurityalliance.org

Security Zoning

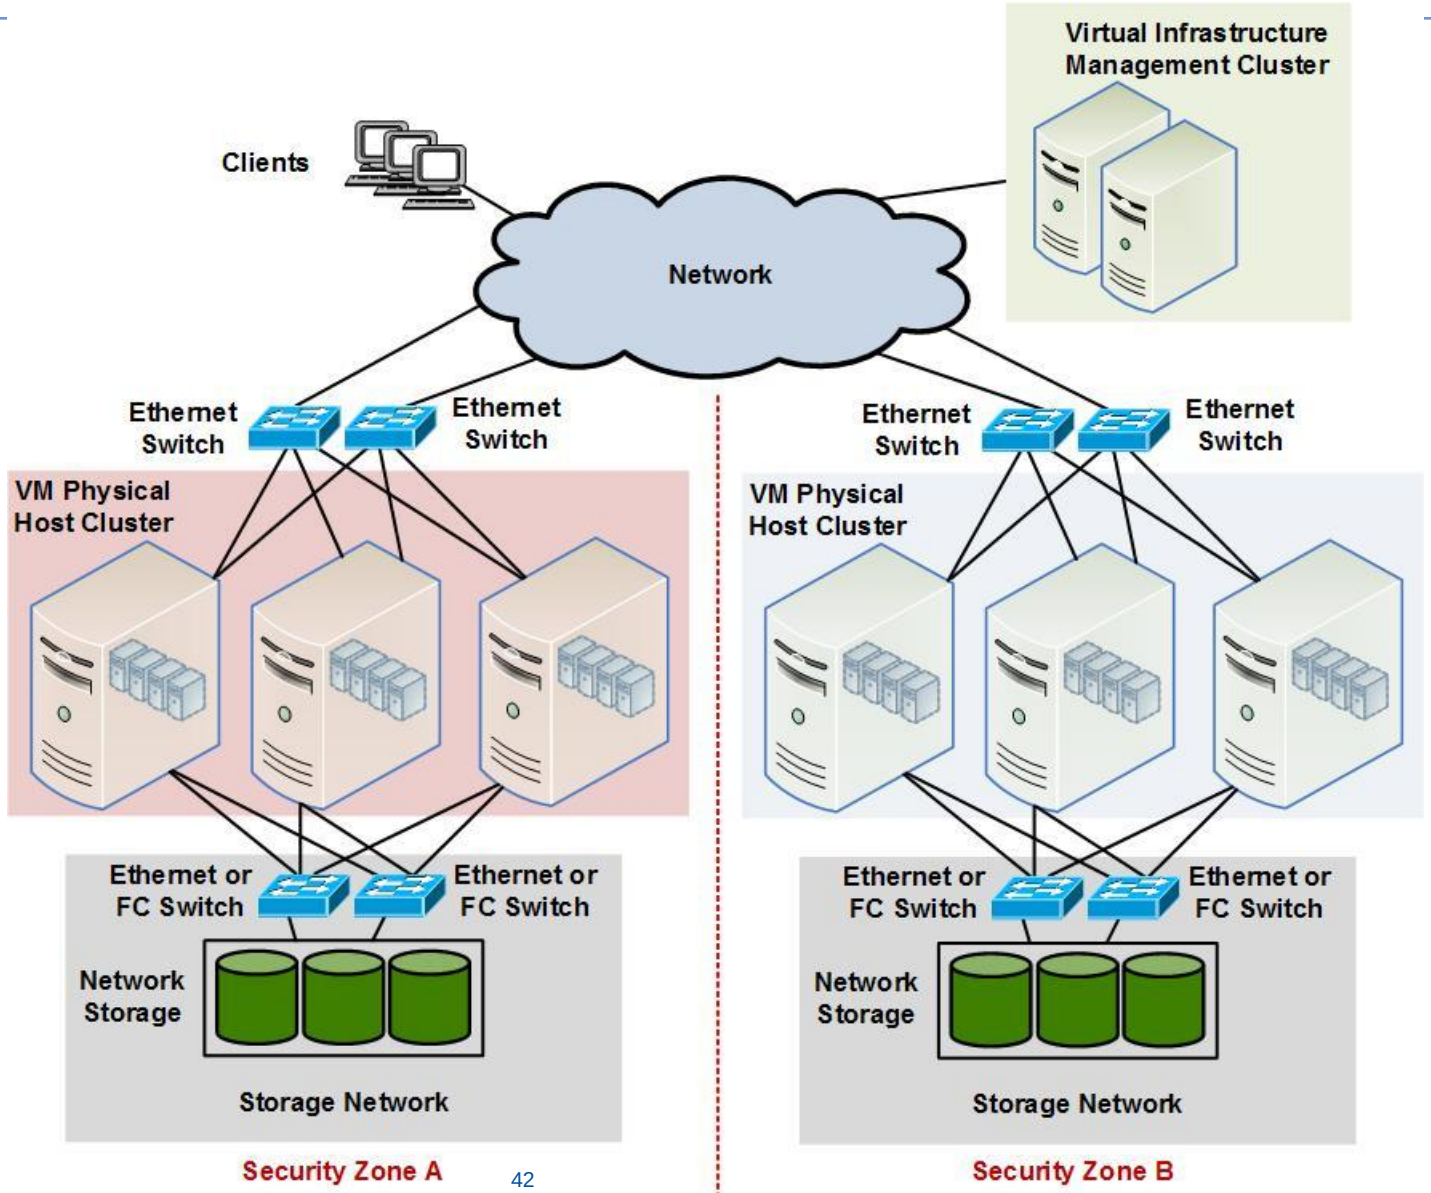

Network Traffic Inspection

- Are you monitoring VM-to-VM traffic?
- Difficult to capture with a physical appliance
  - Requires VLAN trunking to force VM-to-VM traffic to traverse physical infrastructure
  - Only option today for Hyper-V

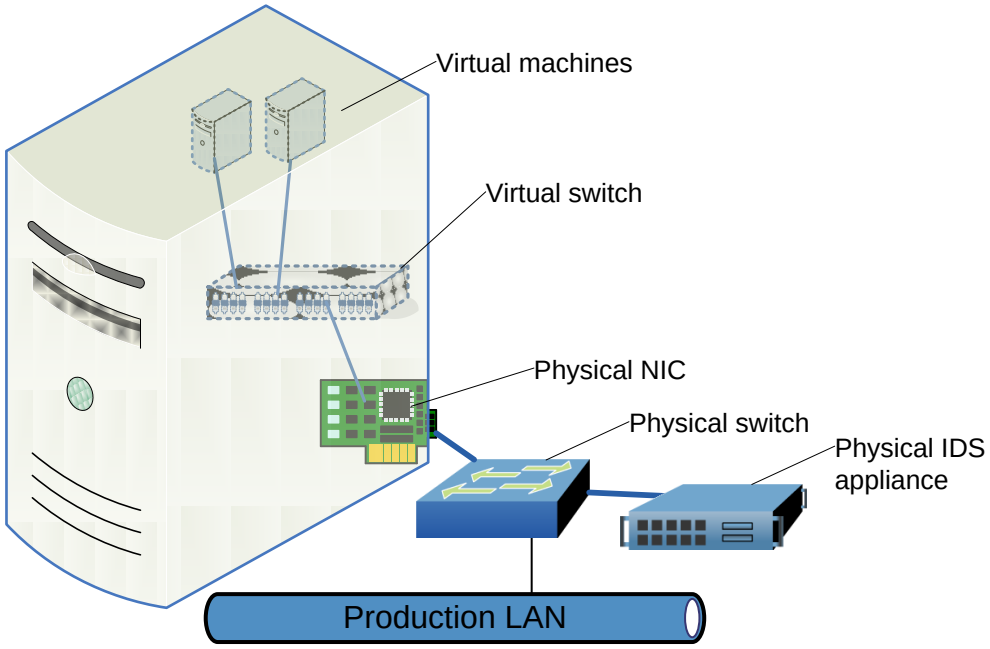

Gartner

Network Traffic Inspection

- VM appliance
- Integrates with existing vSwitch architecture (standard, DVS, Nexus 1000V, Open vSwitch)
- May be required on every physical host

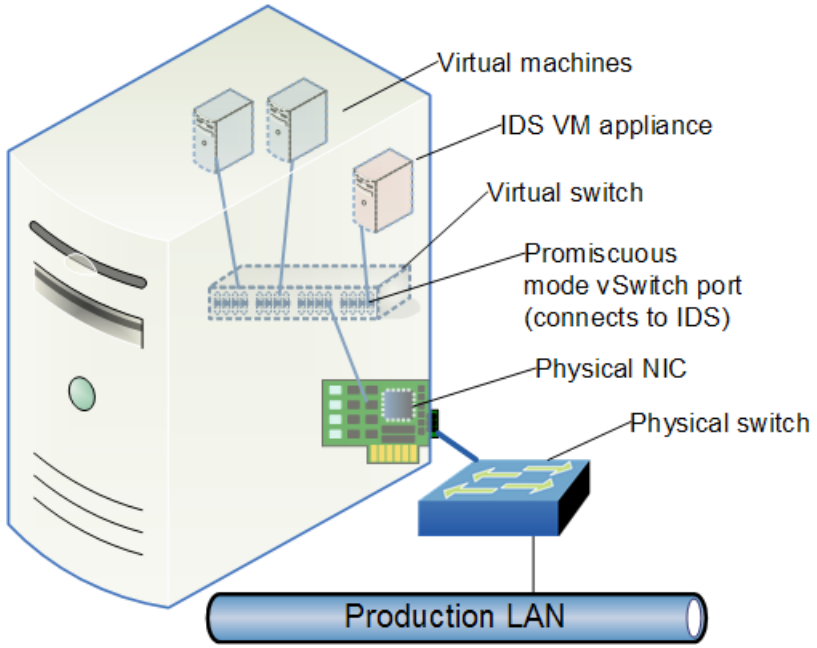

Gartner

Network Traffic Inspection

Introspection

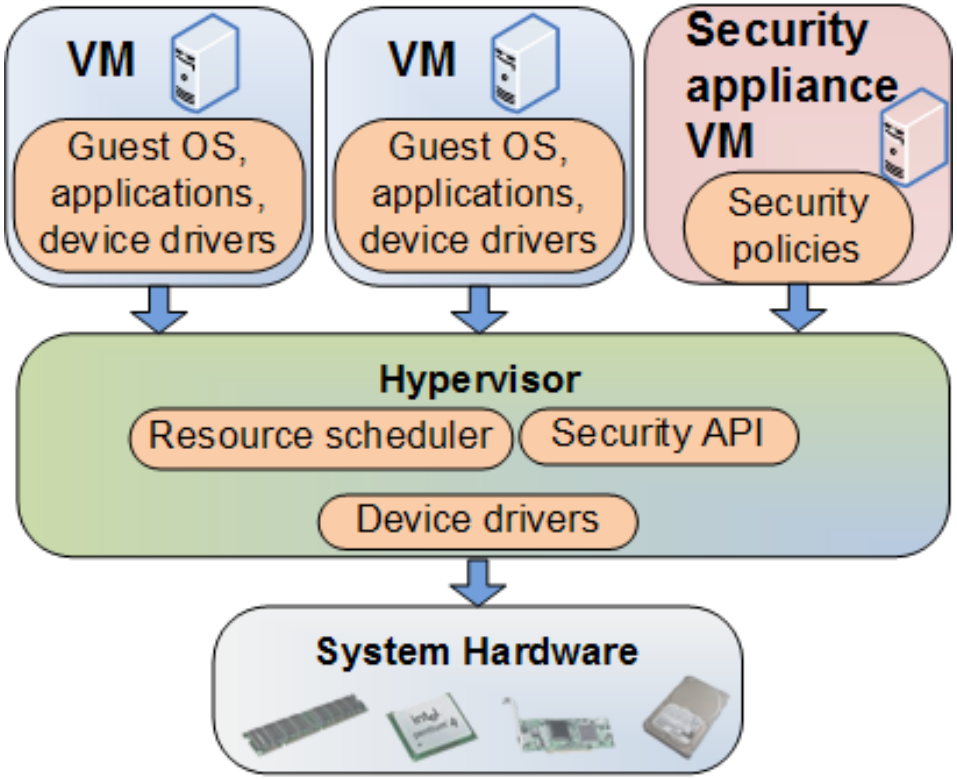

Gartner

Security Best Practices

- Virtual Infrastructure Configuration
  - Use native and third party features to satisfy Gartner's hypervisor security criteria
  - Implement/audit for recommendations in vendor security hardening guides
    - vSphere 4.0: http://www.vmware.com/resources/techresources/10109
    - Microsoft Hyper-V: http://technet.microsoft.com/en-us/library/dd569113.aspx
    - XenServer 5.0.3: http://support.citrix.com/servlet/KbServlet/download/19723-102-82000/user_security_v5update3.pdf

Gartner

Security Best Practices

- Zoning and Physical Isolation
  - Use separate physical clusters to isolate security zones (e.g. DMZ and internal trusted zones)
  - Internal subzones can reside on shared physical infrastructure
  - Subzone isolation provided by:
    - VLANs
    - Dedicated virtual switches/physical NIC ports/back-end storage
  - Some enterprises are mixing zones on shared server infrastructure
    - Using host-based security and planning to add network-based security
    - Dedicating network and storage resources to each zone

Security Best Practices

- Solving the Network Traffic Inspection Dilemma
  - Using host-based security
    - Provides consistent method to monitor, audit, and enforce security, regardless of underlying physical or virtual infrastructure
  - Using network-based security:
    - Provides additive layer
    - Enforces security on VM appliances that do not support installed agents or applications
  - vShield App + vShield Edge
    - **vShield App**: distributed virtual firewall
    - **vShield Edge**: edge firewall for each VDC
      - Stateful inspection, Audit/compliance, site-to-site VPN, load balancing, NAT, and DHCP

Security Best Practices

- “We’re SAS 70 compliant!”
  - Ask provider for audit details
- Not all PCI auditors accept shared physical infrastructure as adequate separation
- Metadata that identifies physical data location and VM runtime location is required
  - RSA/Intel/VMware can identify VM, but not data today
- Virtual data centers will aid VM mobility to cloud; plan for them
- Virtual appliance management belongs to the specialists (network, security admins)
- Key management: Enterprises (not providers) should hold/manage keys

Security Best Practices

- Infrastructure Authority (IA)
  - Central metadata store
  - Maintains dependency maps
    - Compute, memory, network, storage
    - Security and regulatory compliance
    - Data: location, latency requirements, protection, and recovery
    - Environmental (power, cooling)
    - Cost
  - Stores and enforces organizational policy
    - Are security zoning rules checked before live migrating a VM?
    - Do any site restrictions prevent VMs from migrating to different data centers or to public cloud infrastructure?
  - Ensures accurate capacity forecasts

Data Protection

- Status quo is horribly inefficient
  - Agent-based backups are resource-intensive:
    - Reduce VM density
    - Increases infrastructure costs
  - Multiple backup jobs required for each recovery scenario
  - Multiple tools are required to recover data
    - Which tool backed up which server?
    - Which tool is needed for recovery?

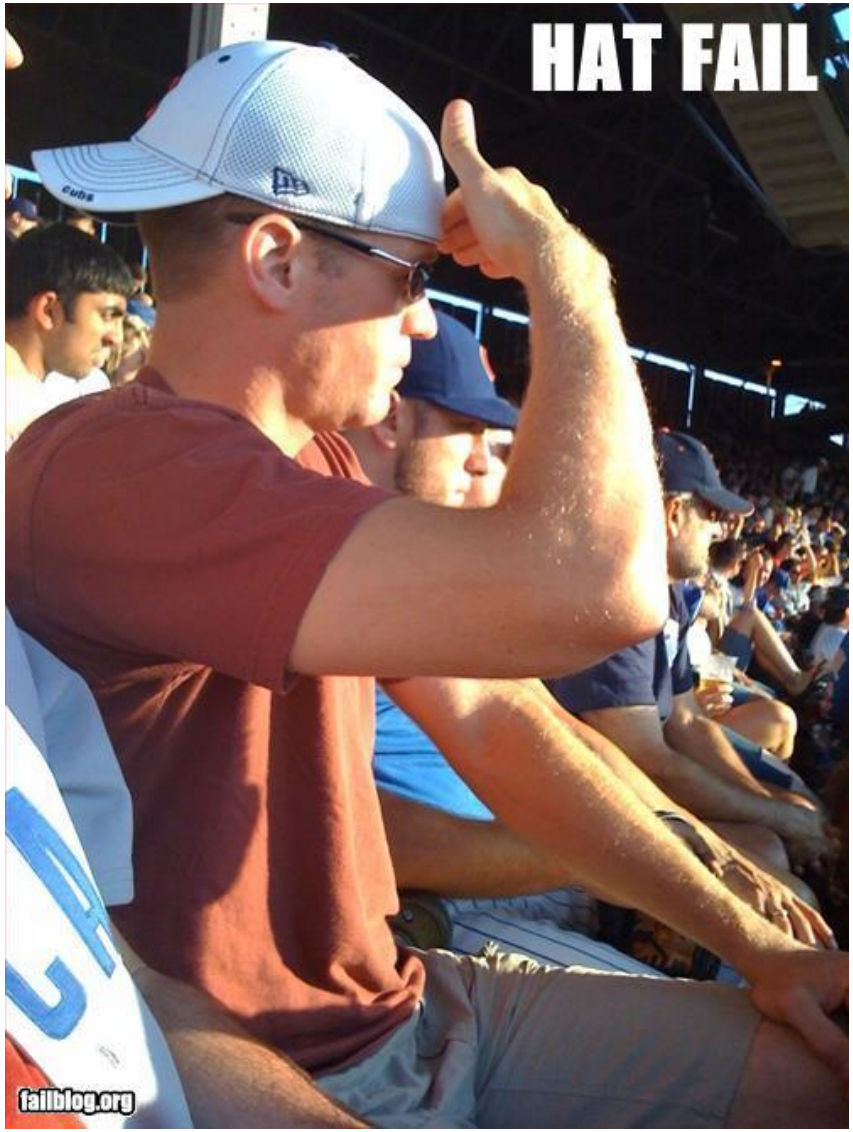

Future Developments

- The Data Protection/Recovery Orchestrator

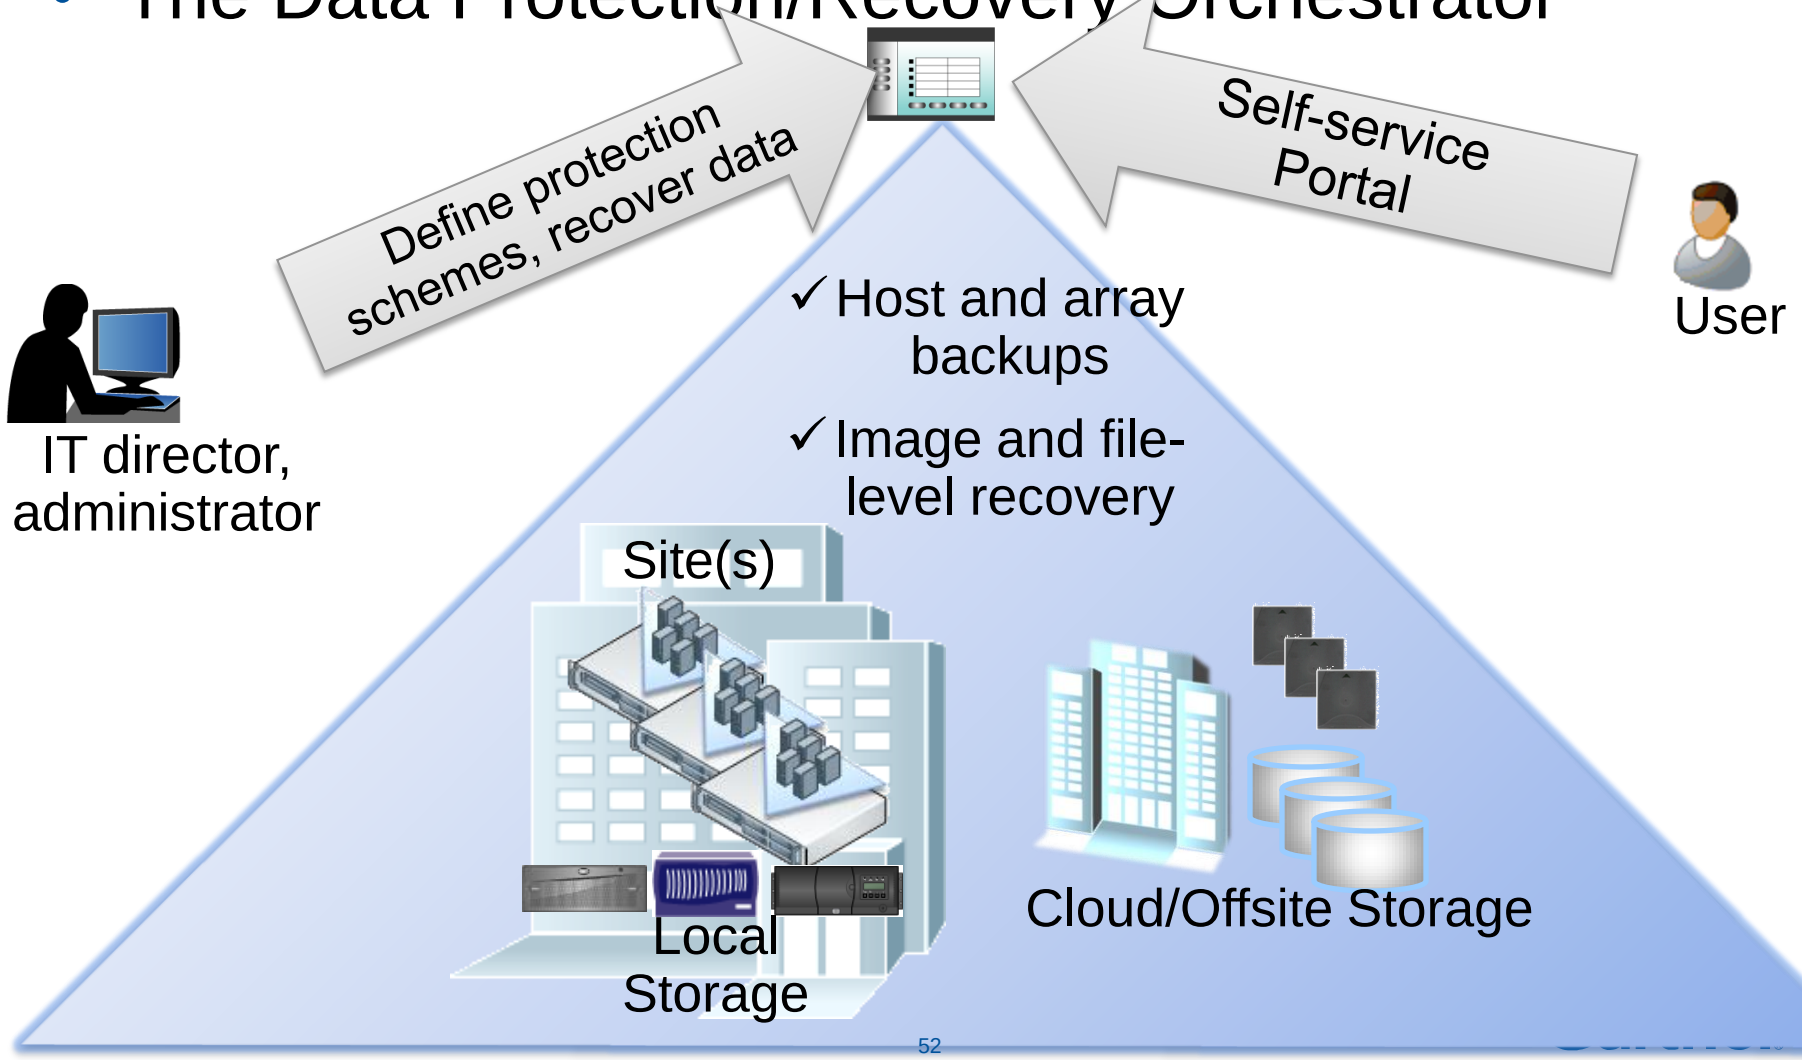

Service Oriented

- Key characteristics:
  - Tier 1 services and applications are virtualized
  - End users are able to provision entire services (e.g., multi-tier application stacks) using a self-service web interface
  - The HIaaS infrastructure is optimized to support multitenancy
  - HIaaS services are indexed in a service catalog
  - A chargeback system is in place

Service Oriented

- Objectives:
  - Optimize for multi-tenancy and virtualize tier 1 apps and services
  - Deploy virtual infrastructure appliances
  - Implement or update a service catalog
  - Automate service provisioning
  - Define HIaaS Standard Models

Cloud Enabled

- Key characteristics:
  - Services are deployed to virtual data centers to maintain security while supporting cloud infrastructure mobility
  - The infrastructure is highly orchestrated and intelligently uses external cloud resources to optimize capacity, business continuity, or to reduce operational costs
  - The internal cloud is fully capable of operating in a hybrid cloud environment

Cloud Enabled

- Objectives:
  - Optimize for cloud portability
  - Implement the IA
  - Integrate cloud infrastructure management with the IA
  - Automate cloud bursting
  - Integrate policy enforcement with the IA

Future Trends

- Infrastructure Authority (IA)
  - Central metadata store
  - Maintains dependency maps
    - Compute, memory, network, storage
    - Security and regulatory compliance
    - Environmental (power, cooling)
  - Stores and enforces organizational policy
    - Are security zoning rules checked before live migrating a VM?
    - Do any site restrictions prevent VMs from migrating to different data centers or to public cloud infrastructure?
  - Ensures accurate capacity forecasts
  - Dynamically expands infrastructure to external cloud

Infrastructure Authority – Practical Steps

- Push for industry standards
- Natural evolution of virtual infrastructure management
  - vCenter, System Center, XenCenter
- Runtime metadata
  - File that identifies all service level requirements
  - Formats:
    - OVF, vmx
  - Vendors must publish a standard schema
- Plug-in architecture
  - Standard XML template
  - Used by providers to "plug-in" to IA
- Standard models (e.g., security, infrastructure, billing)

Summary

- The goal is multitenancy, but we’re not there yet
- Developing organizational standard models for security and service levels helps private cloud transparency
- The virtual data center as an isolation boundary will be broadly accepted over next 2-5 years
- Network access layer appliances and storage infrastructure should be build with virtual data center in mind